

REPORT DOCUMENTATION PAGE

AFRL-SR-AR-TR-03-

data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188/4302). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to provide information if it does not have a valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) 31-08-2003		2. REPORT TYPE		3. DATES COVERED (From - To) 1 April 2000 - 31 May 2003	
4. TITLE AND SUBTITLE Mixed-Initiative Knowledge Base Development				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER F49620-00-1-0072	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Gheorghe Tecuci				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) MSN 4A5, Learning Agents Laboratory, Department of Computer Science, School of Information Technology and Engineering George Mason University 4400 University Dr. Fairfax, VA 22030-4444				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR/NM 4015 Wilson Blvd., Room 713 Arlington, VA 22203-1954				10. SPONSOR/MONITOR'S ACRONYM(S) AFOSR	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT This research has developed an end-to-end mixed-initiative approach to the development of knowledge bases by subject matter experts, with limited assistance from knowledge engineers. In this approach, the complex knowledge engineering activities, traditionally performed by a knowledge engineer and a subject matter expert, are replaced with equivalent ones performed by the subject matter expert and a learning agent, through mixed-initiative reasoning, with limited assistance from the knowledge engineer. In essence, the learning agent helps the subject matter expert to describe a specific problem, to make explicit how he or she solves it, to formalize this reasoning and to explain it to the agent. At the same time, the expert helps the agent to understand this reasoning process, to learn general problem solving tasks and rules from it, and to refine its ontology, thus developing its knowledge base to represent the expertise of the subject matter expert. These methods have been implemented in the Disciple-RKF learning agent that has been successfully used in several courses at the US Army War College, in the context of the center of gravity analysis problem. Experimental results demonstrate that the developed methods simplify the acquisition of knowledge and improve the knowledge base development process.					
15. SUBJECT TERMS artificial intelligence, knowledge acquisition, mixed-initiative reasoning, knowledge engineering, multistrategy machine learning, knowledge base, agent, ontology, problem solving through task reduction, subject matter expert, center of gravity analysis					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 18	19a. NAME OF RESPONSIBLE PERSON Gheorghe Tecuci
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (include area code) 703 993 1722

20031028 181

Mixed-Initiative Knowledge Base Development

AFOSR Grant # F49620-00-1-0072

Final Technical Report

1 April 2000 - 31 May 2003

PI: Prof. Gheorghe Tecuci

MSN 4A5, Learning Agents Laboratory

Computer Science Department, George Mason University

4400 University Dr. Fairfax VA 22030-4444

phone: (703) 993 1722, fax: (703) 993 1710, email: tecuci@gmu.edu

<http://lalab.gmu.edu>

Abstract

This research has developed an end-to-end mixed-initiative approach to the development of knowledge bases by subject matter experts, with limited assistance from knowledge engineers. In this approach, the complex knowledge engineering activities, traditionally performed by a knowledge engineer and a subject matter expert, are replaced with equivalent ones performed by the subject matter expert and a learning agent, through mixed-initiative reasoning, with limited assistance from the knowledge engineer. In essence, the learning agent helps the subject matter expert to describe a specific problem, to make explicit how he or she solves it, to formalize this reasoning and to explain it to the agent. At the same time, the expert helps the agent to understand this reasoning process, to learn general problem solving tasks and rules from it, and to refine its ontology, thus developing its knowledge base to represent the expertise of the subject matter expert. These methods have been implemented in the Disciple-RKF learning agent that has been successfully used in several courses at the US Army War College, in the context of the center of gravity analysis problem. Experimental results demonstrate that the developed methods simplify the acquisition of knowledge and improve the knowledge base development process.

1. Introduction

The objective of this research project was to develop mixed-initiative problem solving, learning, and knowledge reuse methods for knowledge base development, that synergistically integrate the human reasoning and knowledge of a subject matter expert with the automatic reasoning and embedded knowledge engineering methods of an intelligent agent, to take advantage of their complementary knowledge, reasoning styles and computational strengths. Building on previous work on knowledge acquisition through multistrategy learning, this research has developed an end-to-end mixed-initiative approach to the development of knowledge bases by subject matter experts,

with limited assistance from knowledge engineers. This research has synergized with the research done under DARPA's Rapid Knowledge Formation program (2000-2004), leading to the development, experimental use, and transition of a complex knowledge engineering environment, called Disciple-RKF, and its application to the military center of gravity analysis domain.

The next two sections summarize the main results of this research, including the production of two PhD dissertations. After the title of each contribution there are numeric references to the published papers listed in section 4. Section 5 lists the presentations and the demonstrations of the performed research that have been made at the AFOSR PI meetings, at the DARPA RKF meetings, and with other occasions, such as invited talks. However, it does not include the presentations and the demonstrations associated with the conference or workshop papers listed in section 3. Section 6 lists the most significant events, achievements, transitions, and interactions with other organizations, which took place during this research project. Finally, section 6 lists the personnel associated with this research.

2 Summary of the Main Contributions

End-to-End Approach to Mixed-Initiative Knowledge Base Development [12, 15, 30, 36]

This research has elaborated an end-to-end mixed-initiative approach to the development of a knowledge base by a subject matter expert (SME), with limited assistance from a knowledge engineer (KE). In this approach, called Disciple-RKF, the complex knowledge engineering activities, traditionally performed by a knowledge engineer with assistance from a subject matter expert, are replaced with equivalent ones performed by the subject matter expert and a learning agent, through mixed-initiative reasoning, with limited assistance from a knowledge engineer. The expert and the agent collaborate to develop the knowledge base of the agent so that it represents the problem solving expertise of the subject matter expert. The top part of Figure 1 shows the complex knowledge engineering activities that are required to build a knowledge base. The knowledge engineer has to develop a model of the application domain that will make explicit the way the subject matter expert solves problems. Then the knowledge engineer has to develop the object ontology. He or she also needs to define general problem solving rules and to debug them. As shown at the bottom of Figure 1, each such activity is replaced with an equivalent activity which is either performed by the subject matter expert and the agent, or requires limited assistance from the knowledge engineer (KE). The knowledge engineer is still needed to help the subject matter expert to define an initial domain model and to develop an initial object ontology. After that, however, the domain model and the ontology can be extended and refined by the subject matter expert and the agent. The complex activities of defining and debugging problem solving rules require no or very limited support from the knowledge engineer. The subject matter expert can teach the agent how to solve problems, through examples and explanations, and the agent can learn and refine the rules by itself.

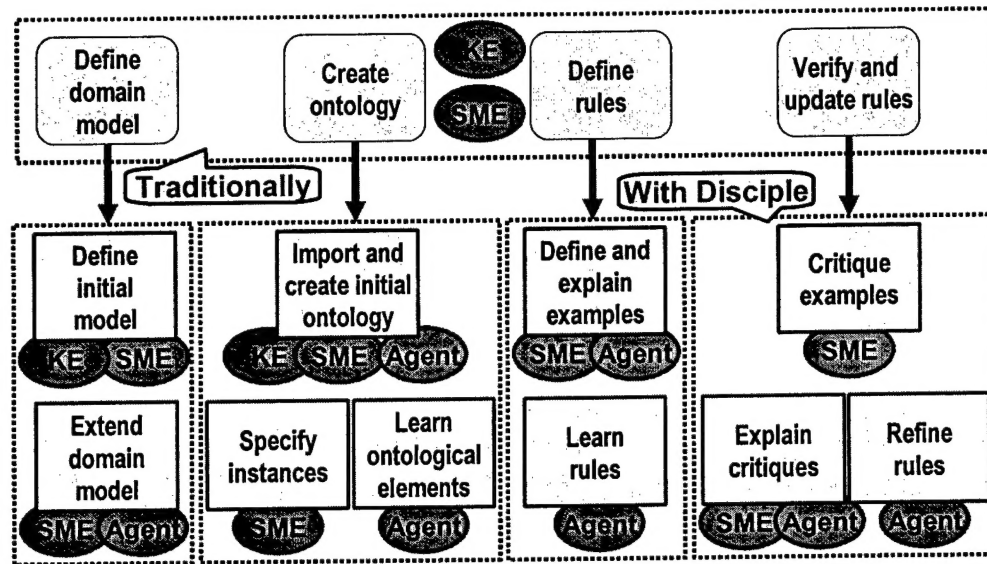


Figure 1: Complex knowledge engineering activities replaced with simpler mixed-initiative activities.

As indicated in Figure 2, The Disciple-RKF approach covers all the phases of knowledge bases and agents development and use, from the initial analysis of an expert's problem solving process, to ontology specification, import, development and learning, to agent teaching, problem solving, rule learning and ontology learning, and finally to knowledge base integration, knowledge base export, and agent use.

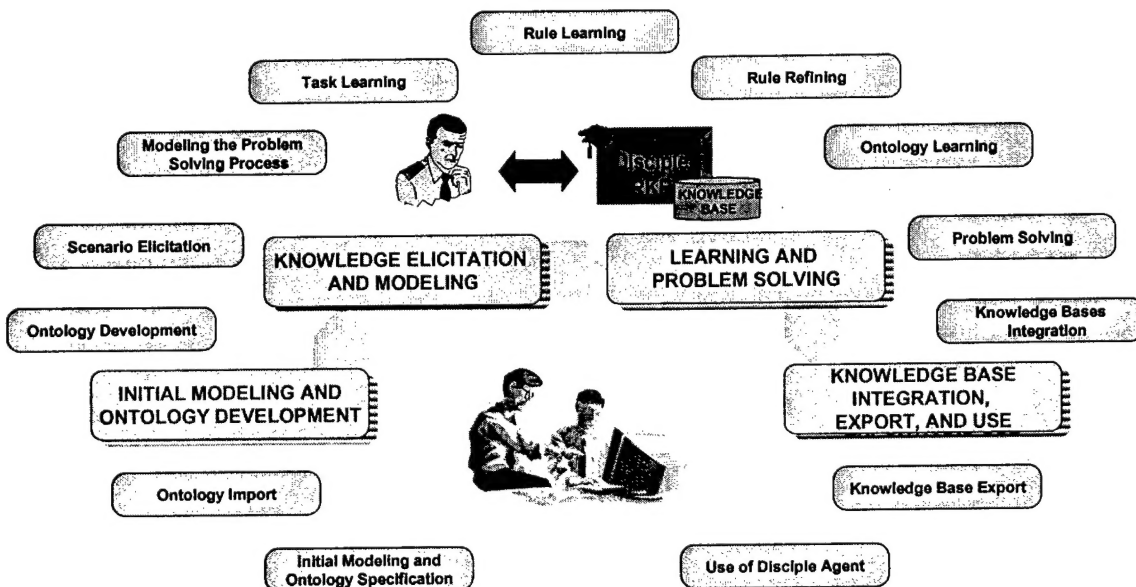


Figure 2: The phases of the Disciple-RKF approach.

Learnable Knowledge Representation [2, 18, 13]

This research has produced a learnable knowledge representation for a knowledge base that supports all the mixed-initiative knowledge base development activities shown in Figures 1 and 2, as summarized below.

The knowledge base is structured into an object ontology that represents the objects from an application domain, and a set of task reduction and composition rules that represents the problem solving expertise of a subject matter expert. This separates very clearly the knowledge which is characteristic to an entire application domain (the object ontology), from the knowledge which is characteristic to a specific subject matter expert (the task reduction rules and the solution composition rules). This clear separation facilitated the development of methods for importing the object ontology from external knowledge repositories, and for learning the rules directly from the subject matter expert.

Previous research on the Disciple approach has introduced the notion of plausible version space to represent a partially learned rule in an evolving representation language. This research project has extended the plausible version space representation, using it to represent all the knowledge pieces from the knowledge base. The object ontology represents the generalization hierarchy for learning. The features of these objects, the problem solving tasks, and the task reduction rules, are learned from specific examples, based on the plausible version space representation.

Another characteristic of the developed knowledge representation, which is very important for knowledge acquisition from a subject matter expert, is that it represents knowledge both at an informal (natural language) level, and at a formal (logic) level. The Disciple-RKF agent uses the formal representation to perform reasoning and the informal one to communicate with the expert.

Concept Learning with an Evolving Representation Language [2]

As mentioned above, the knowledge base of a Disciple agent consists of an object ontology, and a set of problem solving rules expressed with the terms from the ontology. The object ontology, which is also used as the generalization hierarchy for learning, is inherently incomplete, and may not be able to represent the target concept that the agent is attempting to learn. This research has developed a method to represent and learn the best approximation of a target concept in an evolving generalization hierarchy of concepts. The method uses the union, the intersection, and the complementary of the existing concepts to determine an approximation of the concept to be learned. One has introduced the version space of the best approximations of the target concept, which is bounded by a lower approximation and by an upper approximation. It has been proven that, for any approximation of the target concept, there is a better approximation within this version space. A method to compute the lower and upper bounds of this version space has also been developed. Compared to the previous Disciple learning methods, the proposed method has the following advantages: it is based on a more powerful representation language (with union and complementary/negation, in addition to

intersection); it offers a controlled convergence toward the best approximations; and it helps in identifying missing concepts in the ontology. In addition, this new method allowed the refinement of the symbolic approach to uncertainty used previously in Disciple, by distinguishing between several levels of confidence in the classification of new instances. The method is a basic component of several higher level learning methods of Disciple, such as the methods to learn the domains and the ranges of the object features and of the task features, and the methods to learn the applicability conditions of the tasks and of the task reduction rules.

Problem Solving Method Based on Solution Composition Rules

The general problem solving paradigm of the Disciple approach is task reduction, where a general problem solving task is successively reduced to simpler tasks. In the previous versions of the Disciple approach, once the initial task was reduced to elementary tasks, the solutions of these elementary tasks were successively composed, according to a fixed rule, into the solution of the initial task. For instance, in the course of action critiquing domain, the solution of a critiquing task is the union of the solutions of its subtasks [31]. Many problem solving tasks however, such as center of gravity analysis [37], require a complex solution composition process where the composition of the solutions depends of the solutions themselves. To be able to address such problems, the problem solving paradigm of Disciple-RKF was generalized to include a rule-based solution composition process, where simpler solutions are composed into more general ones, based on general solution composition rules. This included the development of a knowledge representation for solution composition rules with plausible version space conditions, the development of structural pattern matching methods for component sub-solutions, and the development of methods for defining and refining solution composition rules. A future research direction is the development of methods for learning general solution composition rules from specific examples provided by the expert.

Scenario Elicitation from a Subject Matter Expert [11, 28]

One has developed a general knowledge elicitation approach that allows the agent to elicit the description of an input problem or scenario directly from the subject matter expert, through a natural dialog. The main idea is to associate elicitation scripts with the general concepts and features from the agent's object ontology. Each such elicitation script describes how the agent should elicit the instances of the corresponding concept, what questions to ask the expert, and how to represent the expert's answer in the knowledge base. This approach allows rapid development of customized scenario elicitation tools by a knowledge engineer, while before such a task needed to be performed by the system developer.

Mixed-Initiative Modeling of the Expert's Problem Solving Process [2, 14]

During the knowledge acquisition experiments performed with subject matter experts at the US Army War College, it has been determined that the single most difficult agent training and knowledge base development activity for an expert is to express in detail his

or her reasoning process, even when using unrestricted natural language. During this process the expert has to explain to the agent how he or she is solving a current problem solving task, using the task reduction paradigm (i.e. what relevant question to ask, what is the answer to that question, and which are the subtasks that result from this answer). To facilitate this complex process, a modeling method has been developed and implemented into a modeling assistant. The modeling assistant uses several heuristic methods, such as: (1) Word completion heuristics based on lexical analysis, reasoning context, and analogy with similar examples. They help the expert to refer to the entities from the agent's ontology; (2) Suggestions on how to complete the example that is currently modeled, such as, what question to ask in order to reduce the current task, which is its answer, and which are the corresponding subtasks. They are based on analogical reasoning, knowledge engineering guidelines, and plausible justifications. (3) Example analyzer which identifies potential problems in the structure of the current task reduction step being modeled, and suggests corrections to the expert.

Mixed-Initiative Methods for Task, Rule, and Ontology Learning [10, 25, 36]

One has developed new and improved methods for task formalization, task learning, rule learning, rule refinement, rule analysis, and ontology learning. All the learning methods are incremental and are based on the plausible version space representation, sharing many components and strategies. The task learning method learns a plausible version space task starting from a task example specified by the expert. First the expert expresses a new problem solving task, in English, as part of the modeling process. Then the expert and the agent collaborate in formalizing the task, with the agent proposing an initial formalization (based on analogy with previously formalized tasks, natural language processing and the formalization context) and the expert improving it. Based on this formalization the agent generates an initial plausible version space task. This partially learned task is then refined based on new task examples encountered during mixed-initiative problem solving. The previously developed explanation generation method (which is part of both the rule learning method and of the rule refinement method) has been improved by automatically generating some plausible "common-sense" explanations, and by simplifying the type of learning hints that the expert can provide to the agent, without reducing the expert's ability to guide the agent. An interactive method for learning general object features from specific features (given by the expert as explanations of problem solving steps) was also developed. One has also developed a method to analyze a learned rule in order to identify the need for additional explanations of the example from which the rule was learned. If some of the variables of the learned rule are not constrained enough, the rule will generate very many solutions, most of which are incorrect. Immediately detecting such unconstrained variables allows the agent to ask questions about the example from which the rule was learned, to identify the missing explanations of the input example, and to generate a better rule.

Learning and Discovery of Ontological Knowledge [5, 6]

Because the agent's ontology is incomplete and is used as the generalization hierarchy for learning, the learned rules generally contain exceptions. Therefore, one has developed a

suite of interactive methods for extending the object ontology with new objects and features in order to eliminate the exceptions of the rules. Some are simpler methods for a subject matter expert, while others are more complex ones, requiring the participation of a knowledge engineer. Each method has four major phases: 1) Discovery of promising ontology refinement candidates; 2) Selection of the candidate; 3) Elicitation and refinement of the ontology based on the selected candidate; and 4) Refinement of the rules based on the refined ontology. The developed methods are based on mixed-initiative human-agent interaction, heuristic analysis of the knowledge base and plausible reasoning. Experiments performed at the US Army War College have confirmed that the rules learned from subject matter experts have a significant number of exceptions that provide valuable information on how the ontology should be extended in order to represent subtle distinctions in the application domain.

Integrated Modeling, Ontology Development, Learning and Problem Solving [2, 3, 4, 34]

An important characteristic of the Disciple-RKF knowledge base development approach is the synergistic integration of modeling, ontology development, learning and problem solving. These activities (which have traditionally been separately treated), are mutually supporting each other. During the initial modeling of the expert's problem solving process a specification of the object ontology is developed. This specification guides the ontology import and development process. Once the object ontology has been developed, the expert can teach the agent how to solve problems by using the previously developed modeling examples. This process involves a mixed-initiative formalization of the informal task reduction steps which result in learning of plausible task reduction rules by the Disciple agent. These rules, in turn, are used during mixed-initiative problem solving and are further refined as a result of their use (no matter whether they were successfully used or not). The refined rules may also accumulate exceptions which are eliminated by further extending the object ontology through ontology elicitation and learning. The learned rules are also used by the modeling assistant mentioned above to support the expert in the modeling process. Therefore, in the Disciple-RKF approach, the modeling, ontology development, learning and problem solving activities mutually support each other.

Mixed-Initiative Ontology Import and Merging Methods [1, 11]

A mixed-initiative ontology import method has been developed. This method consists of three basic phases: 1) Mixed-initiative retrieval of potentially relevant ontological knowledge from an external knowledge repository; 2) Automatic translation of the retrieved ontological knowledge into an intermediate Disciple ontology; and 3) Mixed-initiative import from the intermediate Disciple ontology into the final Disciple ontology. The first phase is done by the subject matter expert and the Disciple agent, the second one is done by the Disciple agent, and the third one is done jointly by the subject matter expert, the Disciple agent, and the knowledge engineer. Based on this ontology import method one has also developed a mixed-initiative method for merging multiple

knowledge bases which have been developed in parallel starting, from a common ontology.

3. PhD Dissertations Completed

A Methodology for Modeling Expert Knowledge that Supports Teaching-Based Development of Agents

Michael Bowman

Ph.D. Dissertation in Information Technology, April 2002

Abstract

In order for artificial intelligence to become truly useful in real-world applications and environments it is necessary to identify, document, and integrate into automated systems the human knowledge that people use to solve their real-world problems. This process has been found to be difficult, and is a critical part of what has become known as the knowledge acquisition bottleneck. The primary contribution of this dissertation is the development and application of a general methodology for modeling and representing an expert's problem-solving knowledge that supports ontology import and development, teaching-based intelligent agent development, and agent-based problem solving. The methodology provides practical guidance to subject matter experts on how to express the way they solve problems using the task reduction paradigm. It identifies the necessary concepts and features to be represented in the ontology; identifies the tasks to be represented in the agent's knowledge base; guides the rule learning and refinement processes; supports natural language generation of solutions and justifications, and is natural and easy to use. The methodology is applicable to a wide variety of domains and has been successfully used in military planning, course of action critiquing, and strategic center of gravity identification problems. This research is part of a much larger effort with the goal to develop an advanced approach to expert knowledge acquisition based on apprenticeship multi-strategy learning, in a mixed-initiative framework.

Available at: <http://wwwlib.umi.com/dissertations/fullcit/3040412>

Modeling and Learning with Incomplete Knowledge

Mihai Boicu

Ph.D. Dissertation in Information Technology, December 2002

Abstract

Our research goal is to allow subject matter experts who have no prior knowledge engineering experience to build intelligent agents that incorporate their problem solving expertise. Such an instructable agent will face the challenge of working with an incomplete and therefore evolving representation space. This thesis presents an integrated

set of methods for knowledge representation, modeling, learning, and problem solving with incomplete knowledge. These methods are synergistically integrated to compensate for each other's weaknesses with their complementary strengths. First, we present a new form of plausible version space rules, based on a new method of concept learning in an incomplete generalization hierarchy of objects. We prove that this method converges almost monotonically to the version space of the best approximations of the target concept. Then, we describe new methods for modeling and learning an expert's problem solving knowledge. These methods have been implemented in the context of the Disciple approach, in three different systems that have been successfully applied to three complex domains, as demonstrated by nine intensive experiments. The experimental results proved that our proposed methods advance the state of the art in instructable agents, giving hope that in the near future typical computer users will be able to train their own personal software assistants.

Available at: <http://wwwlib.umi.com/dissertations/fullcit/3068628>

4. Cumulative List of Publications

- [1] Barbulescu M., Balan G., Boicu M., Tecuci G., "Rapid Development of Large Knowledge Bases," in *Proceedings of the 2003 IEEE International Conference on Systems, Man & Cybernetics*, October 5-8, Washington D.C., 2003.
- [2] Mihai Boicu, "Modeling and Learning with Incomplete Knowledge," *PhD Dissertation in Information Technology*, Learning Agents Laboratory, School of Information Technology and Engineering, George Mason University, 2002.
- [3] Boicu M., Marcu D., Bowman M., and Tecuci G., "A Mixed-Initiative Approach to Teaching Agents to Do Things," in *Proceedings of the Symposium on Learning How to Do Things*, The 2000 AAAI Fall Symposium Series, North Falmouth, Massachusetts, November 3-5, 2000.
- [4] Boicu M. and Tecuci G., "Mixed-Initiative Reasoning for Integrated Domain Modeling, Learning and Problem Solving," in *Proceedings of the Seventeenth National Conference on Artificial Intelligence and the Twelfth Conference on Innovative Application of Artificial Intelligence*, Menlo Park, CA:AAAI Press, 2000.
- [5] Boicu C., Tecuci G., Boicu M., "Mixed-Initiative Exception-Based Learning for Knowledge Base Refinement," in *Proceedings of the Eighteenth National Conference of Artificial Intelligence and the Fourteenth Conference on Innovative Applications of Artificial Intelligence*, AAAI-02/IAAI-02, pp. 947-948, Edmonton, Alberta, Canada, AAAI Press/The MIT Press, 2002.
- [6] Boicu C., Tecuci G., Boicu M., Marcu D., "Improving the Representation Space through Exception-Based Learning," in *Proceedings of the 16th International FLAIRS Conference (FLAIRS-2003)*, Special Track on Machine Learning, May 2003, Key West, Florida. AAAI Press, Menlo Park, CA, 2003.

- [7] Boicu M., Tecuci G., Marcu D., Boicu C., Stanescu B., "Mixed-Initiative Control for Teaching and Learning in Disciple," in *Proceedings of IJCAI-2003 Workshop on Mixed-Initiative Intelligent Systems*, Acapulco, Mexico, August, AAAI Press, Menlo Park, CA, 2003.
- [8] Boicu M., Tecuci G., Marcu D., Bowman M., Shyr P., Ciucu F., and Levcovici C., "Disciple-COA: From Agent Programming to Agent Teaching," in *Proceedings of the Seventeenth International Conference on Machine Learning (ICML)*, Stanford, California 2000, Morgan Kaufman.
- [9] Boicu M., Tecuci G., Marcu D., Stanescu B., Boicu C., Balan C., Barbulescu M. and Hao X., "Disciple-RKF/COG: Agent Teaching by Subject Matter Experts," in *Proceedings of the Eighteenth National Conference of Artificial Intelligence and the Fourteenth Conference on Innovative Applications of Artificial Intelligence*, AAAI-02/IAAI-02, pp. 992-993, Edmonton, Alberta, Canada, AAAI Press/The MIT Press, 2002.
- [10] Boicu M., Tecuci G. and Stanescu B., "Mixed-Initiative Agent Teaching and Learning," in *Proceedings of the 2001 International Conference on Artificial Intelligence*, IC-AI'2001, pp. 122-128. June 25-28, 2001, Monte Carlo Resort, Las Vegas, Nevada, USA, IEEE Computer Society, Los Alamitos, California, 2001.
- [11] Boicu M., Tecuci G., Stanescu B., Balan G.C. and Popovici E., "Ontologies and the Knowledge Acquisition Bottleneck," in *Proceedings of IJCAI-2001 Workshop on Ontologies and Information Sharing*, pp. 9-18. Seattle, Washington, August 2001, AAAI Press, Menlo Park, California, 2001.
- [12] Boicu M., Tecuci G., Stanescu B., Marcu M. and Cascaval C.E., "Automatic Knowledge Acquisition from Subject Matter Experts," in *Proceedings of the Thirteenth International Conference on Tools with Artificial Intelligence (ICTAI)*, pp. 69-78. 7-9 November 2001, Dallas, Texas, IEEE Computer Society, Los Alamitos, California, 2001.
- [13] Boicu M., Tecuci G., Stanescu B., Panait L., Cascaval C., "Learnable Representation for Real World Planning", in *Proceedings of the AAAI-2000 Workshop on Representational Issues for Real-World Planning Systems*, Austin, Texas, 2000.
- [14] Bowman M., "A Methodology for Modeling Expert Knowledge that Supports Teaching-Based Development of Agents," *PhD Dissertation in Information Technology*, School of Information Technology and Engineering, George Mason University, Fairfax Virginia, 2002.
- [15] Bowman M., Lopez A.M. Jr., Donlon J., and Tecuci G., "Teaching Intelligent Agents: Software Design Methodology," *CROSSTALK: The Journal of Defense Software Engineering*, Vol. 14, No.6, pp. 10-14, 2001.

[16] Bowman M., Lopez, A.M. Jr., and Tecuci G., "Ontology Development for Military Applications," *Proceedings of the South Eastern Regional ACM Conference*, Atlanta, GA, May 2001.

[17] Bowman M., Tecuci G., and Boicu M., "A Methodology for Modeling and Representing Expert Knowledge that Supports Teaching-Based Intelligent," in *Proceedings of the Seventeenth National Conference on Artificial Intelligence and the Twelfth Conference on Innovative Application of Artificial Intelligence*, July 2000, Menlo Park, CA: AAAI Press, 2000.

[18] Bowman M., Tecuci G., and Boicu M., "Intelligent Agents in the Command Post," in *Military Review*, March-April 2001, pp. 46-53, US Army CGSG, Fort Leavenworth, Kansas.

[19] Bowman M., Tecuci G., and Boicu M., "Intelligent Agents: Tools for the Command Post and Commander," in *Army Acquisition, Logistics and Technology (ALT)*, pp.34-35, Nov-Dec 2001.

[20] Bowman M., Tecuci G., and Ceruti M., "Application of Disciple to Decision Making in Complex and Constrained Environments," in *Proceedings of the 2001 IEEE Systems, Man, and Cybernetics Conference*, pp. 2932-2940, October 2001.

[21] Cleckner W., Tecuci G., Boicu M., Marcu D., Stanescu B., Boicu C., Comello J., Lopez A., Donlon J., "Teaching Strategic Center of Gravity Analysis Through Learning Agent Technologies," in *Proceedings of the Military Operations Research Symposium*, 6pg, Ft. Leavenworth, June 2002.

[22] Cleckner W., Tecuci G., Comello J., "Instructable Agents for Strategic Center of Gravity Analysis," in *Proceedings of the Army Operations Research Symposium*, AORS-2001, 5pg, October 2001.

[23] Fowler C.W., Cleckner W., Lopez A., Tecuci G., Comello J., "Strategic Center of Gravity Analysis for Military Operations other than War through Learning Agent Technologies," in *Proceedings of the Army Operations Research Symposium*, AORS-2002, October 2002.

[24] Lopez, Jr. A.M., Comello J.J., Bowman M., Tecuci G., Donlon J.J., "Center of Gravity Meets Learning Agent Technology," in *Military Review*, November-December, 2002, US Army CGSG, Fort Leavenworth, Kansas.

[25] Marcu D., Boicu M., and Tecuci G., "User-Agent Interactions in Mixed-Initiative Learning," in *Proceedings of the 14th International FLAIRS Conference (FLAIRS-2001)*, Special Track on Machine Learning, May 2001, Key West, Florida. AAAI Press, Menlo Park, CA, 2001.

- [26] Rezazad H., Tecuci G., "Development of an Intelligent Agent for the Design of Local Area Networks," in *Proceedings of the 6th International Conference on Artificial Intelligence in Design*, Worcester Polytechnic Institute, Worcester, Massachusetts, USA, June 26-29, 2000.
- [27] Shyr P., Tecuci G. and Boicu M., "Evaluation of Mixed-Initiative Knowledge Base Development Methods and Tools," in *Proceedings of IJCAI-2001 Workshop on Empirical Methods in AI*, pp. 47-53, Seattle, Washington, August. AAAI Press, Menlo Park, CA, 2001.
- [28] Stanescu B., Boicu C., Balan G., Barbulescu M., Boicu M., Tecuci G., "Ontologies for Learning Agents: Problems, Solutions and Directions," in *Proceedings of the IJCAI-03 Workshop on Workshop on Ontologies and Distributed Systems*, Acapulco, Mexico, August, AAAI Press, Menlo Park, CA, 2003.
- [29] Tecuci G., Aha D., Boicu M., Cox M., Ferguson G., and Tate A. (eds), *Proceedings of the IJCAI-03 Workshop on Mixed-Initiative Intelligent Systems*, Acapulco, Mexico, August 2003.
- [30] Tecuci G., Boicu M., "Military Applications of the Disciple Learning Agent," Chapter 8 in Jain L, (ed.), *Advanced Information Systems in Defense and Related Applications*, pp. 337-376, Springer Verlag, 2002.
- [31] Tecuci G., Boicu M., Bowman M., and Marcu D., with a commentary by Burke M., "An Innovative Application from the DARPA Knowledge Bases Programs: Rapid Development of a High Performance Knowledge Base for Course of Action Critiquing," in *AI Magazine*, 22, 2, 2001, pp. 43-61. AAAI Press, Menlo Park, California, 2001.
- [32] Tecuci G., Boicu M., Bowman M., Marcu D., Shyr P., and Cascaval C., "An Experiment in Agent Teaching by Subject Matter Experts," in *International Journal of Human-Computer Studies*, pp. 583-610, 53, 2000.
- [33] Tecuci G., Boicu M., Marcu D., "Learning Agents Teachable by Typical Computer Users," in *Proceedings of the AAAI-2000 Workshop on New Research Problems for Machine Learning*, Austin, Texas, 2000.
- [34] Tecuci G., Boicu M., Marcu D., "Toward a Disciple-based Mixed-Initiative Cognitive Assistant," in *Proceedings of IJCAI-2003 Workshop on Mixed-Initiative Intelligent Systems*, Acapulco, Mexico, August, AAAI Press, Menlo Park, CA, 2003.
- [35] Tecuci G., Boicu M., Marcu D., Bowman M., Ciucu F., and Levcovici C., "Rapid Development of a High Performance Knowledge Base for Course of Action Critiquing," in *Proceedings of the Seventeenth National Conference on Artificial Intelligence and the Twelfth Conference on Innovative Application of Artificial Intelligence*, July 2000, Menlo Park, CA: AAAI Press.

[36] Tecuci G., Boicu M., Marcu D., Stanescu B., Boicu C., Comello J., "Training and Using Disciple Agents: A Case Study in the Military Center of Gravity Analysis Domain," in *AI Magazine*, 24, 4, 2002, pp.51-68, AAAI Press, Menlo Park, California, 2002.

[37] Tecuci G., Boicu M., Marcu D., Stanescu B., Boicu C., Comello J., Lopez A., Donlon J., Cleckner W., "Development and Deployment of a Disciple Agent for Center of Gravity Analysis," in *Proceedings of the Eighteenth National Conference of Artificial Intelligence and the Fourteenth Conference on Innovative Applications of Artificial Intelligence*, AAAI-02/IAAI-02, pp. 853-860, Edmonton, Alberta, Canada, AAAI Press/The MIT Press, 2002. (Deployed Application Award).

[38] Tecuci G., Boicu M., Wright K., Lee S.W., Marcu D. and Bowman M., "A Tutoring Based Approach to the Development of Intelligent Agents," Chapter 1 in Teodorescu, H.N., Mlynek, D., Kandel, A. and Zimmermann, H.J. (eds), *Intelligent Systems and Interfaces*, Kluwer Academic Press. 2000.

5. Presentations/Demonstrations at AFOSR and DARPA Meetings, Invited Talks

[39] Mixed-Initiative Knowledge Base Development (presentation and demonstration), AFOSR Program Manager Review, George Mason University, October 3rd, 2000.

[40] Development of the Disciple Approach and its Application to the Center of Gravity Challenge Program (presentation), Army War College, Carlisle, PA December 29th, 2000.

[41] Collaborative Assistant for Rapid Knowledge Formation and Reasoning (presentation and demonstration), DARPA RKF PI Meeting, Tucson, AZ, posted at: http://reliant.teknowledge.com/RKF/meetings/RKF_PI_012401/PI_agenda_012401.htm January 24-26, 2001.

[42] Development of Disciple Agents by Subject Matter Experts (presentation), Joint NIMA - US Navy Acquisition Reform Office, LALAB meeting, George Mason University, February 8th, 2001.

[43] Use of Disciple the COG Course (presentation), COG Course After-Action Review, Army War College, Carlisle, PA, March 9th, 2001.

[44] Mixed-Initiative Knowledge Base Development (presentation), AFOSR Software and Systems Program Review, on CD Rom prepared by AFOSR, Ithaca, NY, April 29-31, 2001.

[45] Overview of Disciple-RKF/COG (presentation), posted at: <http://lalab.gmu.edu/RKF/cog/default.htm>, MAAI Course After-Action Review, Army War College, Carlisle, PA, May 31st, 2001.

- [46] Agent Development in MAAI: Student Research Projects (presentation and demonstration), posted at: <http://lalab.gmu.edu/RKF/cog/default.htm>, MAAI Course After-Action Review, Army War College, Carlisle, PA, May 31st, 2001.
- [47] Agent Development Experiment (presentation), posted at: <http://lalab.gmu.edu/RKF/cog/default.htm>, MAAI Course After-Action Review, Army War College, Carlisle, PA, May 31st, 2001.
- [48] Overview of Disciple (presentation), NIMA-LALAB Meeting, George Mason University, July 7th, 2001.
- [49] Mixed-Initiative Knowledge Base Development: The Disciple Approach (presentation), Technology Fair 2001, Patent and Trademark Office, September 5th, 2001.
- [50] Disciple-RKF/COG: Progress Report (presentation, poster and demonstration), DARPA-RKF PI Meeting, McLean, VA, posted at: http://reliant.teknowledge.com/RKF/meetings/RKF_PI_101701/index.html, October 17-19, 2001.
- [51] Disciple – Center of Gravity (presentation), Army War College, Carlisle, PA November 29th, 2001.
- [52] Disciple-RKF/COG: Research, Education, Practice (presentation, poster and demonstration), DARPA-RKF PI Meeting, San Diego, CA, posted at: http://reliant.teknowledge.com/RKF/meetings/RKF_PI_022702/agenda.htm, February 27 – March 1, 2002.
- [53] Use of Disciple in the COG Course (presentation), Term II, After Action Review, Army War College, Carlisle, PA, March 6th, 2002.
- [54] Use of Disciple in the COG Course (presentation), Term III, After Action Review, Army War College, Carlisle, PA, May 22nd, 2002.
- [55] Use of Disciple the MAAI Course (presentation), Term III, After Action Review, Army War College, Carlisle, PA, May 23rd, 2002.
- [56] Development of Knowledge-based Agents by Subject Matter Experts (presentation), The Ontology Firm, NSA, Bowie, MD, May 29th, 2002.
- [57] Mixed-Initiative Knowledge Base Development: Progress Report (presentation), AFOSR Software and Systems Program Review, on CD Rom prepared by AFOSR, Syracuse, NY, June 3-7, 2002.
- [58] Disciple-RKF (presentation and demonstration), BBN, Arlington, VA, July 9th, 2002.

[59] Disciple – Center of Gravity: Use in the AWC Courses (presentation), US Army War College, Carlisle, PA, September 25th, 2002.

[60] Disciple-RKF/COG (presentation, poster, and demonstration), DARPA-RKF PI Meeting, Hilton Head Island, SC, posted at: http://reliant.tekknowledge.com/RKF/meetings/RKF_PI_2002-11-13/index.html, November 12-15, 2002.

[61] Disciple – Center of Gravity: Present and Future (presentation), US Army War College, Carlisle, PA, March 4th, 2003.

[62] Use of Disciple in 319JW - Term II (presentation), After-Action Review Meeting for the COG course, US Army War College, Carlisle, PA, March 14th, 2003.

[63] Disciple-COG (presentation and demo), US Army War College and Center for Army Analysis Exchange Meeting, US Army War College, Carlisle, PA, April 17th, 2003.

[64] Development and Use of Intelligent Decision-making Assistants: the Disciple Approach (presentation), Invited talk, Decision Analysis for Strategic Leaders course, US Army War College, Carlisle PA, May 5th, 2003.

[65] Disciple-RKF/COG: Progress Report (presentation, poster and demonstration), DARPA-RKF PI Meeting, San Diego, posted at: http://reliant.tekknowledge.com/RKF/meetings/RKF_PI_2003-05-13/index.html, CA, May 12-14 2003.

[66] Disciple-COG: Experiments with Learning Agents Technology (presentation and demonstration), Joint After-Action Review Meeting for the COG and MAAI Courses, US Army War College, Carlisle, PA, May 29th, 2003.

[67] Mixed-Initiative Knowledge Base Development: Progress Report (presentation and demo), AFOSR Software and Systems Program Review, Syracuse, NY, June 2-6, 2003.

[68] Personal Learning Assistants (presentation and demonstration), Air Force Research Laboratory, Rome, NY, June 2nd, 2003.

[69] Collaborative Assistant for Rapid Knowledge Formation and Reasoning, (presentation and demonstrations), RKF-PBA Interchange Day, Arlington, VA, posted at: http://reliant.tekknowledge.com/RKF/meetings/062703_RKF_PBA_interchange/index.html, June 27th, 2003.

6. Significant Events, Transitions, and Interactions

Transition to the DARPA's Rapid Knowledge Formation Program

The basic research performed with AFOSR support was applied to the development of several modules of the Disciple-RKF, the Collaborative Assistant for Rapid Knowledge Formation and Reasoning, as part of the DARPA's Rapid Knowledge Formation program. Disciple-RKF includes the following modules: the scenario elicitation module, several ontology browsers and editors (tree-based ontology browsers for objects and features, graph-based hierarchical browsers for objects and features, a graph-based association browser, an object viewer and an object editor, a feature viewer and a feature editor), the modeling assistant, the task formalization and learning module, the rule learning module, the rule refinement module, the exception-handling module, the interactive problem solving module, the autonomous problem solving module, the solution composition editor, the ontology elicitation and learning module, the ontology import module, the ontology merging module, the knowledge base management module, and the reports generation module.

The development of Disciple-RKF, and the experimentations performed at the US Army War College (see below), led to experimental results supporting the claims that the developed methods make knowledge acquisition less time consuming, less error-prone and more efficient, improving the knowledge base development process.

Transition to the US Army War College

This research has also benefited from the collaboration with the Center for Strategic Leadership and the Department of Military Strategy, Planning, and Operations of the US Army War College. The goal of the collaboration was to apply Disciple to the Center of Gravity (COG) domain, and to use it in several US Army War College courses. The US Army War College supported this research effort with scenarios, subject matter expertise, experimentation and evaluation. As a result, since Winter 2001, successive versions of Disciple-RKF were used, on a regular basis, in two joint warfighting courses, "319jw Case Studies in Center of Gravity Analysis," and "589jw Military Applications of Artificial Intelligence."

For its use in 319jw Case Studies in Center of Gravity Analysis, Disciple-RKF was taught to incorporate instructor's expertise in center of gravity analysis. In the course Disciple guides the students to specify the relevant aspects of an assigned war scenario. Then it identifies and tests the strategic center of gravity candidates for that scenario, and generates a center of gravity analysis report. The students study and critique the solutions generated by Disciple and finalize the report.

A total of 55 high-ranking US officers from all the military services and the national reserve, and international fellows have attended 6 sections of this course, using and evaluating the developed Disciple agents for center of gravity analysis. All the students from Spring 2003, who have used the most recent version of Disciple, have agreed or

strongly agreed with the following statements: *"The use of Disciple is an assignment that is well suited to the course's learning objectives"* and *"Disciple should be used in future versions of this course"* (the evaluation questionnaire also contained "neutral", "disagree" and "strongly disagree" options). *The significance of this application is that it shows that the Disciple approach can be used to build practical agents for complex real-world problems.*

In the "589jw Military Applications of Artificial Intelligence" course the students teach personal Disciple agents their own reasoning in center of gravity analysis, participating in unique experiments of agent training and knowledge base development. The three sessions of this course (taught between 2001 and 2003) have been attended by a total of 38 US and international officers from all the branches of the military. At the end of these experiments 10 of them strongly agreed, 22 agreed, 7 were neutral and only one disagreed with the statement *"I think that a subject matter expert can use Disciple to build an agent, with limited assistance from a knowledge engineer."* This is an excellent result, given the difficulty of the addressed problem. In the 2003 experiment, the officers have used historic, current and hypothetical scenarios, with both state and non-state actors, to teach personal Disciple agents how to test center of gravity candidates based on the concepts of critical capabilities, critical requirements, and critical vulnerabilities, which have been recently adopted into the joint military doctrine. In addition, for the first time, this experiment also included the merging of the knowledge bases developed in parallel by the officers. *This is the first time that such an experiment has been performed, demonstrating Disciple's capability for rapid and parallel development of knowledge bases by subject matter experts, with limited assistance from knowledge engineers.*

US Army War College Recognitions

US Army War College has provided several recognitions of this research. Gheorghe Tecuci has been appointed Chair of Artificial Intelligence in the Science and Technology Division, Center for Strategic Leadership, US Army War College, for two consecutive years, 2001-2002 and 2002-2003. Cristina Boicu, Mihai Boicu, Dorin Marcu, Bogdan Stanescu, and Gheorghe Tecuci have received the Centennial Coin of the US Army War College from Major General Robert R. Ivany, in 2002. Gheorghe Tecuci has received the coin of the Center from Strategic Leadership in 2001. Mihai Boicu, Dorin Marcu, Bogdan Stanescu, and Cristina Boicu have received a Certificate of Appreciation from Prof. Douglas Campbell, Director of the Center for Strategic Leadership, in 2001, and from Major General Robert R. Ivany, the commandant of the US Army War College, in 2002.

Deployed Application Award

The paper "Development and Deployment of a Disciple Agent for Center of Gravity Analysis," by Tecuci G., Boicu M., Marcu D., Stanescu B., Boicu C., Comello J., Lopez A., Donlon J., Cleckner W., has received the Deployed Application Award from the American Association for Artificial Intelligence, at the Fourteenth Annual Conference on Innovative Applications of Artificial Intelligence, IAAI-2002, Edmonton, Alberta, Canada, 30 July – 1 August, 2002.

IJCAI-03 Workshop on Mixed Initiative Intelligent Systems

Gheorghe Tecuci (chair) and Mihai Boicu, co-organized an international workshop on Mixed Initiative Intelligent Systems, as part of the 2003 International Joint Conference of Artificial Intelligence, Acapulco, Mexico (see <http://lalab.gmu.edu/MIIS/default.htm>).

Artificial Intelligence Magazine Papers

The Fall 2002 issue of the AI Magazine "highlighting the best work from IAAI-02" (the Innovative Applications of Artificial Intelligence conference) contains the joint George Mason University – US Army War College paper: Tecuci G., Boicu M., Marcu D., Stanescu B., Boicu C., Comello J., Training and Using Disciple Agents: A Case Study in the Military Center of Gravity Analysis Domain. Also, the Summer 2001 issue of the AI Magazine contains the GMU paper "An Innovative Application from the DARPA Knowledge Bases Program: Rapid Development of a Course of Action Critiquer," by Gheorghe Tecuci, Mihai Boicu, Michael Bowman and Dorin Marcu, with a commentary by Murray Burke, the program manager of the DARPA's HPKB and RKF programs.

Disciple-RKF Demonstrated at AAI-02

The Disciple system was demonstrated at the Eighteenth National Conference on Artificial Intelligence at Edmonton, Canada, July 30- August 1, 2002 as part of the AAI-02 Intelligent Systems Demonstrations.

Disciple Approach Seen as Revolutionary Software Design Methodology

The June 2001 special issue on "Software Design Methodologies" of the CrossTalk journal includes an article on the Disciple approach which is introduced by Kevin Richins, the publisher of this issue, as: "This is quite a revolutionary approach that may cause you to rethink traditional defense software development methodologies."

SIGNAL Magazine on the LALAB/GMU Research

The February 2001 issue of the Armed Forces Communications-Electronics Association (AFCEA)'s SIGNAL magazine (<http://afcea.org/>) includes the article "Intelligent Agents Get Smarter" by Henry S. Kenyon, that describes the research done in the Learning Agents Laboratory of George Mason University.

7. Personnel Associated with the Research Effort

Faculty: Gheorghe Tecuci, Mihai Boicu
Students: Bogdan Stanescu, Michael Bowman, Cristina Boicu (Cascaval), Dorin Marcu, Marcel Barbulescu, Gabriel Balan, Elena Popovici, Ping Shyr, James Donlon.